

O
F
S
D

Ad Hoc Networks:
Headline 2000 Communications
Analysis

W.D. Blair and A.B. Reynolds

DSTO-TR-1317

DISTRIBUTION STATEMENT A
Approved for Public Release
Distribution Unlimited

20021119 139



Ad Hoc Networks: Headline 2000 Communications Analysis

W.D. Blair and A.B. Reynolds

Communications Division
Electronics and Surveillance Research Laboratory

DSTO-TR-1317

ABSTRACT

This report examines unit location data and terrain from Headline 2000 to investigate communications networks within the manoeuvre units. Within the Enhanced Combat Force timeframe, such units would be supported by a Tactical Data Distribution Subsystem (TDDS). This is envisaged as an *ad hoc* network changing its topology as units manoeuvre across the battlespace. The report describes approaches that can be applied to Headline 2000 and other data to characterise the nature of the network topology, both statically and as it changes over time, in order to explore the impact on the TDDS network capacity and hence to assist in elucidating consequent requirements of a candidate TDDS.

RELEASE LIMITATION

Approved for public release

AQ F03-02-0360

Published by

*DSTO Electronics and Surveillance Research Laboratory
PO Box 1500
Edinburgh South Australia 5111 Australia*

Telephone: (08) 8259 5555

Fax: (08) 8259 6567

© Commonwealth of Australia 2002

AR-012-358

June 2002

APPROVED FOR PUBLIC RELEASE

Ad Hoc Networks: Headline 2000 Communications Analysis

Executive Summary

In the next 15 years or so land forces at brigade and below will employ a communications system known as a Tactical Data Distribution Sub-system (TDDS) as the primary means to meet the data (especially real-time) demands of the force in the future. The TDDS is envisaged as an *ad hoc* network changing its topology as units manoeuvre across the battlespace and is a capability that does not currently exist within the Australian Defence Force.

This report examines unit location data and terrain from recent wargames to consider possible communications networks within the land manoeuvre formation. The report describes approaches that can be applied to Headline 2000 and other data to characterise the nature of the network topology, both statically and as it changes over time, in order to explore the impact on the TDDS network capacity and hence assist in elucidating consequent requirements of a candidate TDDS. The work also contributes to DSTO's long range research task on *ad hoc* networks (LRR 01/090) including studies being conducted at the University of Adelaide and the University of South Australia.

In contrast to our use of Australian specific wargame data, much of the literature on *ad hoc* networks uses variations on randomly moving nodes and flat terrain or uses foreign doctrine/terrain. The results of this report and ongoing effort in this field are expected to contribute to Australian Defence Organisation outcomes and capability by contributing to project definition studies for TDDS within Joint Project 2072 – Battlespace Communications System (Land/Air).

Authors

W.D. Blair

Communications Division

William Blair is a Senior Research Scientist in Network Architectures Group of the Defence Science and Technology Organisation. He has been involved in research into the application of commercial communications techniques to military strategic and tactical communications since leaving the Australian Army.

A.B. Reynolds

Communications Division

Alfred Reynolds is a Research Engineer in Network Architectures Group of the Defence Science and Technology Organisation's (DSTO) Communications Division having joined DSTO after completing a Systems Engineering Degree at the Australian National University. He is involved in research into tactical communication systems.

Contents

GLOSSARY

1. INTRODUCTION.....	1
2. BACKGROUND.....	1
2.1 Headline 2000	1
2.2 Graph Theory Background.....	3
3. AD HOC NETWORKS	4
3.1 Ad hoc Network Resource Sharing	5
3.2 Characterisation of Ad hoc Networks	7
3.2.1 Static Characterisation	7
3.2.2 Dynamic Characterisation.....	8
4. HEADLINE 2000 – GRAPHICAL ANALYSIS	9
4.1 Initial Examination	9
4.2 Initial Findings – Static Characterisation	12
4.2.1 Connectedness.....	12
4.2.2 Degree and Density of the System.	15
4.2.3 Phases Required	16
4.2.4 Path Length Factors	19
4.3 Initial Findings – Dynamic Characterisation.....	19
4.3.1 Outage	19
4.3.2 Churn.....	22
5. POTENTIAL FOR FUTURE WORK.....	22
6. CONCLUSIONS.....	23
7. RECOMMENDATIONS.....	24
8. REFERENCES	25
APPENDIX A : HEADLINE 2000 DATA LOCATIONS	26
APPENDIX B : GRAPH THEORY MEASURES	27
B.1 Degree	27
B.2 Density.....	27
B.3 Churn	28
APPENDIX C : STDMA AS A GRAPH PROBLEM.....	30
APPENDIX D : LINKING TO UNIVERSITY OF ADELAIDE TRAFFIC MODEL.....	32

GLOSSARY

TDDS	Tactical Data Distribution Sub-system
ECF	Enhanced Capability Force
RaNAT	Radio Network Analysis Tool
ADFA	Australian Defence Force Academy
BCSS	Battlespace Command Support System
LOD	Land Operations Division, DSTO
CSMA	Carrier Sense Multiple Access
MACAW	Mobile <i>Ad Hoc</i> Network
STDMA	Spatial reuse Time Division Multiple Access
TDMA	Time Division Multiple Access, STDMA
FDMA	Frequency Division Multiple Access
CDMA	Code Division Multiple Access
UAV	Unattended Aerial Vehicles

1. Introduction

Within the timeframe of the Enhanced Capability Force (ECF, i.e. the land forces of 15 years hence), a communications system known as a Tactical Data Distribution Sub-system (TDDS) would support units at brigade and below. The TDDS is that element of the overarching Battlespace Communications System (Land) that provides the primary means to meet the data (especially real-time) demands of the force. Further details are available in [1]. The TDDS is envisaged as an *ad hoc* network changing its topology as units manoeuvre across the battlespace.

The Headline series of wargames are a part of the Army Experimental Framework of studies into the ECF. The experiment depicts a brigade sized force operating using envisaged future weapon systems and adapted tactics. Military participants dictate the deployments with some computer wargame assistance to ensure realistic speeds of movement etc. The deployments are influenced by terrain and operational contact with the enemy; we believe this results in a quite realistic force laydown.

This report examines unit location data and terrain from Headline 2000 to consider possible communications networks within the land manoeuvre formation. The report describes approaches that can be applied to Headline 2000 and other datasets to characterise the nature of the network topology, both statically and as it changes over time, in order to explore the impact on the TDDS network capacity and hence to assist in elucidating consequent requirements of a candidate TDDS. The work also contributes to DSTO's longrange research task on *ad hoc* networks (LRR 01/090) including studies being conducted at the University of Adelaide and the University of South Australia.

In contrast to our use of Australian specific wargame data, much of the literature on *ad hoc* networks uses variations on randomly moving nodes and flat terrain. More realistic land tactical network topologies are tested in work sourced from the Swedish Defence Research Agency (such as [2] and others), but these are based on Swedish tactics/terrain and are still machine generated.

2. Background

2.1 Headline 2000

(From [3]):

Headline 2000 was an experiment conducted within the Army Experimental Framework looking at the development of Command and Control (C2) aspects and Combat Service Support (CSS) aspects in Manoeuvre Operations in the Littoral Environment (MOLE). It was in the form of a two level Command Post Exercise (CPX) with the higher level simulating the Brigade (Equivalent) HQ and the lower element simulating sub-ordinate units (Battle

Groups (BG)). Headline 2000 was conducted with a scenario involving an armoured force and findings from it will be used in the design of 1 Bde in the Enhanced Capability Force (ECF). The control for the experiment was 1 Bde enhanced in an evolutionary manner to the expected capability of 2015 but operating under current doctrine. Two variations on the experimental force (EXFOR1) were part of the experiment. The first, EXFOR1 Team 1, had conventional headquarters, the second, EXFOR1 Team 2, had reduced headquarters with the staffing elements of all the headquarters consolidated in a combat support element (CSE) not collocated with any of the commanders. An opposing force (OPFOR) was structured to oppose each Bde force and was understood that it would remain constant throughout the experiment. The terrain of the experiment was the fictitious country Kamaria.

The units being emulated in Headline 2000 were mostly ground vehicle mounted however a number of helicopters and unattended aerial vehicles (UAVs) were included. The area of operations in the experiment was approximately 300 km by 300 km and is characterised as having significant vertical relief. While the terrain of Kamaria is based on south-east Australia, the orientation has been shifted and the scale has been distorted slightly. In Australian terms, the operation was conducted in the area from Wagga Wagga towards Yass.

Headline 2000 was conducted as a command post exercise with the Janus wargame used as the adjudicator of results of combat and speed of movement. At approximately 3 second intervals, the Janus program logged the location of any entities that had moved during the period. This log was made available to the authors as a credible source of ECF tactical laydowns. Digital elevation data has been developed for use within Janus and this has also been made available. (Access to the Headline 2000 raw data is via Land Operations Division (LOD) and discussed in Appendix A.) This report documents results from one problem (run 92501) within the Headline 2000 series. It is felt that this run provided a good average case from the set of available data. Further analysis can be carried out on the other runs in due course using the approach described in this report.

The availability of digital elevation data permits the calculation of radio path loss between nominated sites using part of the functionality of Communications Division's Radio Network Analysis Tool (RaNAT). The path loss calculator within the RaNAT uses an algorithm provided by University College (ADFA) to calculate the path loss in dB for a given frequency and antenna height – this algorithm is the same as that obtained for use in Army's Battlespace Command Support System (BCSS). The path losses in this paper were calculated assuming a frequency of 400 MHz and an antenna height of 3m (above ground level). These values were used as it was felt they are representative values for a typical, military TDDS network. The RaNAT parses unit location data (developed from the Janus movement logs) and produces, for each Janus time period, a matrix of radio link losses. With knowledge of the maximum acceptable

path loss, this link loss matrix can be translated into link state matrix where a '1' indicates a link is possible and a '0' indicates an unacceptable path loss. (For the purposes of this analysis an indicative value of path loss of 130dB is used – this value is readily achievable with modern radio devices of the configuration required for a TDDS.)

2.2 Graph Theory Background.

In analysing networks such as a TDDS supporting manoeuvre units, it is useful to consider the network as a (mathematical) graph. Since the authors come from a communications background, communications terms are generally used in describing the graphs rather than those from mathematical graph theory. Thus, we state that a graph comprises two classes of entities: nodes and links. A link interconnects two nodes (typically, but not necessarily, different nodes). Links can be directed – indicating that a flow from node to node can occur in the specified direction. Such directed links can be uni-directional or bi-directional. Other graphs (called undirected) do not seek to specify directions of flows or other relationships between the two connected nodes. In a general graph a node might be connected to another node via multiple links. For our TDDS example, the nodes represent the radio devices and the links represent potentially available radio paths (i.e. that which have a path loss lower than the maximum acceptable path loss). Bi-directional links are required for realistic TDDS links. Also, at most one TDDS link can exist between a particular pair of nodes and links from a node to itself are not relevant.

The following are a number of concepts from graph theory that will be used in this report:

- A graph is 'connected' (or form a 'clique') if there is at least one path, i.e. a series of links and nodes, between every pair of nodes in the graph. In communications terms this means that data can be sent from any node to any other node, perhaps via relaying through intermediate nodes.
- Nodes that are directly connected via a single link are 'neighbours'. In communications terms, it is possible for neighbouring nodes to receive a direct transmission from each other.
- The 'degree' of a node is the number neighbours it has. In communications terms, it is the number of other nodes that can potentially receive a direct transmission from a reference node.
- For a connected graph there is a path between every pair of nodes. If one assumes that every link in the graph is of the same length or cost, then the path between two given nodes with the smallest number of links is the 'shortest path' between them. In communications terms the length (i.e. number of links) of this shortest path represents the minimum number of times a message has to be transmitted whilst being relayed from one node to the other (sometimes known as 'hops'). The hop 'diameter' (or sometimes 'girth') of a graph is the maximum length of all shortest paths.

A graph can be represented in a number of ways. A graphical representation of a simple four node network is shown in Figure 1.




Figure 1 - Graphical Representation

Such a representation is useful for simple graphs and human interaction, but poorly suited to computer processing. Another more appropriate approach is via link matrices. The matrix representation corresponding to Figure 1 is given in Figure 2.

	A	B	C	D
A		1		
B	1		1	1
C		1		
D		1		

Figure 2 - Matrix Representation

Each row and each column represent the connectivity of a particular node. A '1' is present in the intersecting cell if the two nodes are neighbours (i.e. have a direct link). The blank cells contain an implicit '0' indicating the lack of a link.

An important note to make is that the concept of links is suggestive of each link out of a node being independent. However, when the graph is representing a radio network that can operate in a broadcast mode, all neighbours can potentially receive the transmitted message from a node simultaneously. This clearly has an impact on the traffic carrying capacity of the network, but also has important implications that will be discussed in respect to possible self-interference in the network. In our literature search, the authors did not find much evidence of graph theoretic work using topologies that allowed broadcasting – it is an area that may bear further investigation.

3. *Ad hoc* Networks

A 'mobile *ad hoc* network' (sometimes known as a MANET) is an autonomous system of mobile nodes connected by wireless links – the union of which form an arbitrary

graph. The nodes are free to move randomly and organize themselves arbitrarily; thus, the network's wireless topology may change rapidly and unpredictably. Such a network may operate in a standalone fashion, or may be connected to other networks. (This definition is based on [4]).

An *ad hoc* network is fundamentally made up of peer to peer connections. There is no need for dedicated relay base stations that nodes connect into. The concept of an *ad hoc* network as presented in this report requires that all nodes be able to relay messages for nodes not directly connected. The concept of 'autonomy' reflects that there is no centralised monitoring or control of the network topology or traffic routing. These decisions are made by the individual nodes with only perhaps some dialog with their immediate neighbours.

3.1 *Ad hoc* Network Resource Sharing

When a collection of nodes operates as a wireless network, there is a requirement for them to share the radio bandwidth (and hence capacity) in some fashion. The capacity of the resource can be shared amongst the transmitters in a number of ways: on the basis of time (i.e. time can be split into time-slots and these allocated to competing uses), frequency (the total radio bandwidth can be split into sub-bands) or via spread spectrum codes. In much of the literature, the term 'phase' is used for this generic resource and this term will be used in our report. Depending on the network topology, such phases can be re-used multiple times across the network without causing interference. In assigning the phase, the first requirement is that a transmitting station not access the phase (transmit) simultaneously with a transmission from its intended receiving station – this would be a case of primary interference. The next challenge in developing the sharing discipline is that of secondary interference (this incorporates the more difficult issue known as the 'hidden station problem' – see [5]). In this case two transmitters (in the more challenging case 'hidden' from each other, i.e. neither station can receive transmissions directly from the other) seek to transmit to a common destination station simultaneously and interfere at the receiver. If these problems can be solved, then the network can, without fear of disruption, re-use phases within the total network and also allow simultaneous transmissions (increasing network capacity).

In implementing this phase allocation, there are two fundamental approaches:

- Demand driven (or contention based) approaches see nodes seeking access to the resource when there is traffic to be sent. Contention based schemes generally perform well under light load (especially where there are no hidden stations), and respond to diversity in traffic loads across the network, but generally do not offer guaranteed access. Examples are simple carrier sense multiple access (CSMA) or more sophisticated protocols such as MACAW.
- Other approaches assign the phase in a more permanent fashion. Such contention free schemes are often preferred when real-time data is involved as more deterministic network access delay (and hence end-to-end delay) can be

provided than with contention based approaches. Time Division Multiple Access (and its *ad hoc* network variant, Spatial reuse Time Division Multiple Access, STDMA) is a typical example and the basis for much of the later discussion of this report.

The STDMA allocation problem is functionally the same as allocation of sub-bands in a (Spatial reuse) Frequency Division Multiple Access (FDMA) network or codes in a direct sequence spread spectrum Code Division Multiple Access (CDMA) network [6]. Examining the STDMA allocation is useful as it can also be argued that the capacity available under STDMA is indicative of the capacity that would be available under a well behaved contention based system when traffic is moderately busy and evenly distributed.

The authors generally refer to STDMA in the discussion in this report as many implemented networks follow this approach, but this should not be construed that any system envisaged for the Australian Defence Force would limit itself to STDMA.

The allocation of the phases can be carried out either on a node (broadcasting transmitter) or link (transmitter-receiver pair) basis. The differing performance of link allocation vs node allocation under different traffic loads is examined in [7]. This concludes that for higher traffic loads link assignment is preferred while node assignment is preferred for lower traffic levels. It further concludes that, from the point of view of end to end delay, as the network topology comes closer to being (fully) meshed the traffic load crossover point (the point where better method of allocation changes from node allocation to link allocation) also increases. That is, as network topology approaches being meshed the traffic load of the network has to be higher to make the link assignment better¹. It should be noted that the paper finds it axiomatic that node assignment is better when traffic is broadcast (i.e. all messages are intended for all stations). The authors contend that a large portion of situational awareness traffic will be of this nature. Most of the literature discovered by the authors employs node assignment and we will limit further discussion in this report to node allocation.

Furthermore, there are two fundamentally different ways to carry out the node allocation approach.

- The first has a fixed number of phases (typically equal or greater than the number of nodes), each node has a guaranteed phase and the challenge is to allocate the other phases such that no interference occurs (such as [8]).
- The second seeks to optimise to the minimum number of unique phases required and allocate a single phase to each node whilst maintaining non-interference (such as [6] and [9]). The non-interference requirements (both primary and secondary) lead the problem to be one that has been described as

¹ The concept of a "meshed" network (or "fully" meshed as some communicators would term it) and its metrics are discussed in para 3.2.1.

an L(1,1) graph colouring problem (specifically in [10], though their description is strictly incorrect as discussed in Appendix B).

At this stage our studies have focussed on the second approach and it is discussed in more detail in Appendix B. Nevertheless, the first approach has some innate benefits (every node has a guaranteed phase) that may simplify network management in a real network. We intend returning to that approach in the future.

3.2 Characterisation of *Ad hoc* Networks

3.2.1 Static Characterisation

The following characterisations of the *ad hoc* network graph have been examined because they contribute to our understanding of network performance implications:

- **Connectedness.** In any real network deployment, terrain will impact on the ability of nodes to connect to their peers. Depending on the nature of the terrain and the location of the tactical forces (and hence nodes), the subsequent *ad hoc* network may not be a single connected graph. This means that some users (or connected groups of users) will not have any path into the remainder of the network. The analysis of connectedness scopes the problem of interconnecting these segmented subnets via a second tier of communications supplementing the terrestrial radio network being modelled.
- **Degree.** The degree of any node in the graph will have an impact on the efficiency of broadcast transmissions (i.e. the number of stations that can simultaneously receive the transmission). The maximum degree in a graph (and the statistics of degree across the nodes in the graph) will be a factor in determining the share of the network resource that a node can expect to be able to use. The degree is a measure of either competing nodes in a contention schemes or will influence the number of time slots that are required to implement an STDMA plan.
- **Density.** (Entitled 'connectivity' in [7]). This is a measure of how close a graph is to being (fully) meshed. The density of a graph is defined as the number of links divided by the number of links possible. In an undirected graph, the number of links possible is given by $N(N-1)/2$ (where N is the number of nodes in the graph). A meshed network has a density of 1; a totally unconnected network. (i.e. a group of individual isolated nodes) has a density of 0². A dense network offers a large number of alternate paths between any two nodes and hence is potentially more survivable.
- **Phases Required.** In the absence of a simple analytical approach to calculate the number of phases (e.g. STDMA time slots) using degree and density

² An alternative way of thinking of this is: a meshed undirected network has all nodes being neighbours to all other nodes thus all nodes would have a degree of $N-1$; a fully disconnected undirected network has all nodes with degree of 0. Density can be calculated as *average degree*/ $N-1$.

characteristics (which we sought to develop and discussed in Appendix B), a 'greedy' algorithm was applied to determine a full allocation plan leading to a reasonable estimate for this parameter. Assuming each node is only allocated one phase to transmit, then the larger the number of phases the proportionately smaller will be the share for each station. Thus if we know the bit rate that the radio channel can sustain, the number of phases shows the capacity that a single node can transmit.

- **Path Length Factors.** There are two measures that fall into this category. The first is the diameter of the graph. The diameter is the longest path (hops) to reach a node from any other node using the shortest path between the two nodes. The diameter gives a value for the maximum number of hops that a packet could travel within the graph. When given a latency introduced per link it provides an upper bound on the possible latency in the network. The second measure is the average path length. This is the average of the shortest path between every node. The average path length provides the average number of hops that a packet will traverse between any two arbitrary nodes in the graph. It provides a good indicator as to the average number of times a packet will be relayed. By comparing the diameter and the average path length it is possible to get an understanding as to the general shape of the graph. Two similar values suggest a largely homogenous network, whereas two distinctly different values point towards a graph that contains a larger cluster and then a few outliers.

3.2.2 Dynamic Characterisation

There are three key elements that allow an examination of the dynamics of the network.

- **Consideration of Acceptable Outage Times.** The interaction between terrain and tactical deployments is examined in a static sense by the connectedness issues above. Dynamic analysis of outage times looks at the statistical characterisation of the separation from the network over the course of the exercise. User input can then be sought to assess operational acceptability of the outage statistics. Examination will help in identification of special nodes whose mission will typically take them out of range of the main network. For a real network, alternative means can be scoped and scaled to service nodes for which outages exceed acceptable limits.
- **Churn.** The concept of churn is introduced in this paper. The extent of change in the network topology will have an impact on the performance of the network (user capacity being impacted through the network self-organising). Churn is a weighted linear combination of normalised measures of how many neighbours have arrived or departed from a node's neighbourhood. The strict definition is:

$$\text{churn} = \alpha \frac{\text{new neighbours}}{\text{current number of neighbours}} + \beta \frac{\text{departed neighbours}}{\text{number of neighbours in previous time step}}$$

(for this work we have set α and β to 0.5). High churn and churn observed at a large number of nodes will require greater network configuration action.

- **Turbulence of the Phase Plan.** This metric has not been pursued to date but would consider how much the phase plan would need to be modified in the face of the topological change. It is envisaged that there would be some correlation between these changes and the analytically derived churn factor, but it remains for future work.

4. Headline 2000 – Graphical Analysis

4.1 Initial Examination

The first aspect of initial examination of the Headline 2000 data focussed on connectedness of ground forces to the UAV fleet. While the UAV fleet was deployed as a surveillance asset, we wished to explore their possible use as communications relay platforms. Our initial thoughts were that a surveillance mission profile would be inimical to operation as a communications relay however this has not been borne out by the results to date. Available documentation stated that the particular UAV would operate at an elevation of 5,000 m above ground level. On the basis of that elevation and the flight plan extracted from Janus, Figure 3 shows the proportion of ground based elements that can sustain a radio path with one or more of the (typically fourteen) deployed UAVs.³

³ Note that whereas the maximum acceptable path loss used for the analysis of the terrestrial was 130 dB, this analysis used 120 dB – this allows for the possibility that the UAV relay is challenged for power/antenna gain etc.




Figure 3 - Connectedness of nodes to UAVs

If the UAVs remain as ordinary peers in the network, they would become central points for the entire *ad hoc* network. Instead, we conceive the UAV to offer a 'link of last resort' to reduce traffic load on the UAV payload, yet bring unconnected clusters or individual users into the network. Thus the UAVs become a communications tier above the terrestrial network. We removed UAV entities from further network analysis to examine the performance of the terrestrial *ad hoc* network.

Initial examination of network characteristics around the helicopter elements was indicative that they should be removed from the *ad hoc* network and linked back via a separate means. They seemed to be the cause of excessive churn and spikes of the maximum degree in the graph. Nevertheless, closer examination led to rejection of that idea as other entities contributed in a similar fashion.

Because of the manner in which the data was extracted from the Janus wargame, and constraints of the game itself, there may be some limitations in the node location data from Headline 2000.

- The location information was extracted from the Janus Movement Log rather than via periodic location reporting. There is a possibility of stationary players in the wargame remaining 'invisible' because they do not move or move only part way through the experiment.
- Janus operates within limited geographic bounds (a 300 km by 300 km 'play box'). Entities that are not within this area do not exist as far as Janus is concerned. As entities depart the play box, there is nothing in the movement log

to indicate they have left the game. Accordingly, entities leaving the game remain in our location data at the point at which they were last reported moving to. These two anomalies result in those arriving in the game appearing arbitrarily on the terrain even though they would notionally have been part of the network for some time prior. Those leaving the game remain extant at the edge of the play box. When multiple entities leave at the same point (for instance along a road) they appear to congregate at a single location with consequent excellent connectivity. An example of this effect is shown in Figure 4 and Figure 5.




Figure 4 - Convoy (top centre) Approaching Edge of Play Box




Figure 5 - Convoy (single cross top centre) has left the Play Box

Figure 4 shows a small convoy approaching the edge of the play box (running from top left to bottom right in the top right corner of the image). Figure 5 shows the situation after the convoy has left the game, with all of the exiting entities now appearing in the position of the yellow cross.

The extent of this problem was considered and we chose to discard a 'startup' period and 'shutdown' period. Timestep 300 was selected to be the first for analysis and the stop point was chosen to be 3299 (giving 3000 timesteps).

4.2 Initial Findings – Static Characterisation

4.2.1 Connectedness

There were up to approximately 180 operational entities (nodes) in Headline 2000. While the need for mutual fire support keep some elements of the force clustered, these clusters themselves are dispersed over a large (200km by 200km) area. In addition to the absolute distances, the dispersion is across a complex terrain. Accordingly one would not expect all entities to be always in a single connected terrestrial radio network and this has been shown to be correct. The sizes (i.e. the number of nodes) of the connected networks (referred to as subnets in this report) are shown in Figure 6. Note there are generally two large subnets and several much smaller subnets. The two large subnets occasionally join. Analysis of the operational entities in each subnet shows is that this largely reflects the two major force elements deployed. The 3D graph is interpreted as follows:

- in different colours and going into the page (to the left) are the various subnets;

- the vertical axis shows how many of the operational entities are members of each particular subnet;
- the time steps are the (approximately 3 second apart) time increments of the Janus reports through the exercise.




Figure 6 - Net sizes over the life of the simulation

Two examples for different time steps are given in Figure 7 and Figure 8 to show the distribution of numbers in the connected nets.




Figure 7 - Net sizes for time step 2800

Figure 7 shows a typical time step of the dataset. It consists of two main subnets of roughly equal size and then a number of significantly smaller subnets, disconnected from the majority of the entities. It shows the relatively small scale of the 'disconnected' problem that may require solution via a second tier network (eg UAV relay). There are a small number of subnets that would require interconnection via another radio means and a relatively small number of individuals that may need to be connected.




Figure 8 - Net sizes for time step 1100

Figure 8 highlights an interesting effect. This figure contains one large subnet consisting of almost every entity. This has been caused by a small number of units (typically 1 or 2) creating a 'bridge' of connectivity between two disconnected subnets. This effect typically lasts for only a few time steps, and is caused by units appearing at (and then disappearing from) the tops of hills and other advantageous locations. This behaviour would be very disruptive to routing algorithms as it greatly alters the number of entities in a terrestrially connected subnet and it also produces a bottleneck in communications along that bridging entity. There may be some benefit in seeking to constrain the ability for major subnets to join using the base tier, but remain operating via a second tier linking the subnets.

4.2.2 Degree and Density of the System.

The next two measures calculated are the degree and density of the system. These two values were explored in an attempt to find a simple relationship between the degree/density of the graph and the number of phases required. The results of the analysis did not provide any obvious insights to a relationship between the measures, the details of this analysis can be found in Appendix B. A potential reason for the lack of a simple relationship could be that density is a network wide measure, whereas what is needed is density around the key nodes that are driving the number of phases.

There are difficulties in determining this local value, the resulting complexity is similar in effort to simply performing the phase allocation algorithm itself.

4.2.3 Phases Required

Figure 9 displays the number of phases required to meet the phase plan over the period of interest as compared with the maximum degree. The peak value of around 100 would lead to each station having guaranteed around 3 kbps transmit (given an indicative channel capacity of 300 kbps).




Figure 9 - Number of Time Slots in Phase Plan over time

There is an apparent relationship between the maximum degree of the graph and the number of phases required in the graph. A simple empirical formula of the relationship for this exercise is that the number of phases required is approximately 1.2 times the maximum degree node in the graph. Examination will be required to see whether this empirical calculation would apply more generally to all networks arising from such tactical deployments.

The phase allocation algorithm takes a very simple, single time step view to the allocation. At each time step the algorithm starts allocating from the maximum degree node, until all the nodes in the graph have a phase. While this is a relatively simple calculation to perform, it does not produce good dynamic behaviour, as a node will

most likely be allocated a different phase at the next time step. This causes 'churn' in the allocation plan that results in higher communication overheads. The results are useful for a static analysis, and to provide an upper bound on the phase requirements, but not for analysing the dynamics of the allocation.

The spikes in Figure 9 are of particular concern for two reasons:

- The transmit data capacity available per host is inversely proportional to the number of phases (more phases mean less capacity).
- A large change in the number phases requires significant communication amongst the nodes to settle on a new phase allocation plan.

The driver for the spikes appears to be a 'spurious' node appearing at an advantageous location (as discussed in 4.2.1) and thereby allowing a large number of possible radio links to the unit.

Further analysis was performed to determine the effect and characteristics of these spikes. There are two fundamental subnets present in the simulation (as can be seen in Figure 6). The periods around the phase spikes were examined in more detail. It was found that only one node has an unusually high degree when the number of required phases spikes up (this node was not the same between each spike event however). The nodes that were causing these spikes were all aerial units, mainly helicopters. Note that during our analysis helicopter units were not modelled with correct altitude information, they were locked to 3 m above the ground as all the other units were (this is due to the logs not having altitude information).

The single high degree node was then examined in more detail. If it has more links (a higher degree) than the number of nodes in a single subnet then it must join multiple subnets. However, the larger of the two subnets contains in the order of 110 nodes and the maximum degree node had below 100 links. A more detailed analysis was therefore performed to determine the identity of nodes being joined.




Figure 10 - Effect on clustering during degree spikes

Figure 10 shows the percentage of links that the maximum degree node has into each fundamental subnet. For each of the time steps with a phase spike the links of the node in question were examined. The figure shows that the high degree node is joining the two subnets together. It also shows that the node has roughly the same number of links into each subnet.

The effect of joining the subnets is two fold. Firstly it reduces the capacity available per node, as the fixed channel capacity must now be assigned amongst more phases. It also creates signalling traffic as the nodes reorganise the phase plan to accommodate the new, highly connected node. What is perhaps even worse is the fact that these spikes are very short lived; the disturbance they induce is quickly removed thereby removing the need for the phase plan changes in the first place.

This analysis indicates that there may be value in maintaining separation of the major functional subnets (retaining linkage through the second tier). This would reduce the number of phases required for the subnet and therefore increase the per node capacity available. The enforced separation of subnets would also remove the resource wastage to the momentary (rather than permanent) joining of the subnets.

4.2.4 Path Length Factors




Figure 11 – Average Path Length and Diameter over time

As with Figure 9, spikes can be seen in the diameter of the graphs being examined. These spikes are caused by the same effect in the phase allocation routine, the joining of two large subnets via a bridging node.

During those periods when the network does not merge into the single majority net, the average path length (shown as the purple graph in Figure 11) is reasonably stable at approximately two hops. Each broadcast data packet (i.e. data that is sent to every node) will therefore have to be transmitted on average two times, meaning that each node will have to commit half of its capacity allocation to broadcasting relayed packets.

4.3 Initial Findings – Dynamic Characterisation

4.3.1 Outage

Statistics for station outage, i.e. not belonging to a network of at least 40 nodes, are shown in Figure 12. The threshold of 40 nodes was derived from Figure 6. This number was selected specifically for this exercise to identify stations not part of a major subnet

and hence likely not connected to a major manoeuvre unit. The graph shows columns (merged into an area because of the large number of entities) representing data points for each operational entity showing the percentage of the time that the entity was in the exercise and was connected. The entities are sorted by percent time connected. The graph is interpreted thus: if operational requirements demand that entities must be connected for 70% of the time at a minimum, then (for this experiment) 177 of 210 (i.e. 84%) of the entities will be supported by the terrestrial *ad hoc* network (this is shown more clearly in Figure 13). The remainder will need to be supported by some other means.




Figure 12 - Outage Times for Each Entity

Figure 13 examines the statistics of the period of outage. The left vertical axis applies to the line graph. The continuous line replicates the information in Figure 12. The right vertical axis is in minutes and applies to the columns. The columns show the maximum contiguous period of outage (peak outage duration) for each operational entity. This is worth examining as, for instance, an entity could be disconnected 50% of the time, but if these are spread in many short periods, then this might still be acceptable. Clearly, total capacity for receiving information would be impacted. However, providing the outages were short enough and evenly spread, such a node might in effect receive

every second report if these were high priority and sent periodically. The arrows on the figure diagrammatically denote areas of concern. Entities are of concern if:

- peak outage duration exceeds a threshold (columns intrude above the horizontal line), or
- total outage time exceeds a threshold (entity is to the right of the vertical line).




Figure 13 - The relationship between total time disconnected and peak outage duration

Figure 13 provides an interesting result. There is a strong correlation between the total time disconnected and the peak outage duration of the disconnection. The findings suggest that the entities that are of concern because of peak outage duration are generally also a problem for total connectivity. This greatly simplifies the process of identifying the type of entities that would need alternative communication assets.




Figure 14 - Detailed peak outage duration statistics

Figure 14 provides a more detailed view of the peak outage duration. Other than those for which total connectivity time is of concern, no units exceed our nominal threshold. What is more interesting however is that no units violate even a 5 minute threshold if they achieved an acceptable total connectivity. This suggests that as long as certain total connectivity times are honoured then peak outage duration can be safely ignored. Operational input (from Defence users) will be required to validate thresholds of concern for peak outage duration and total disconnection time.

4.3.2 Churn

This metric will become more important when consideration is made of the overheads of network maintenance. Details of the churn calculations can be found in Appendix B.

5. Potential for Future Work

This report discusses the background metrics and provides results for one run of the Headline 2000 experiment. Analyses of other runs and other exercises remain (see para 2.1).

The extension of traditional graph theory to address broadcast concepts may be a fertile area for future work (see para 2.2).

No work has been done in this study on the alternative STDMA approach, described in para 3.1, where a fixed number of phases are allocated such that no interference occurs.

The ability of real networks to calculate and disseminate STDMA plans, including the relationship with the churn metric, has not been explored here. In addition, as cited in para 3.2.2, the ability of particular phase allocation methods to cope with dynamic network behaviour and consequent turbulence of the phase plan has not been examined.

The inter-relation between the metrics produced in this document and the traffic model being produced by the long-range research task into *ad hoc* networks (LRR 01/090) has yet to be explored in any detail. Appendix D presents a brief discussion on the limited examination conducted to date. This brief analysis shows that the TDDS capacities indicated by our capacity determining metrics are in the appropriate order to carry situational awareness traffic. More detailed analysis is required.

There may be value exploring whether there is a simple metric to describe the 'strength' of a network given the possible relationship between density and survivability.

As mentioned in paras 4.2.1 and 4.2.3, intermittent bridging of two large networks of entities caused problems with phase allocation. Mechanisms to prevent intermittent joining of relatively static subnets have not been explored in this study.

6. Conclusions

Key conclusions are:

- There are some concerns about the underlying location data used for this analysis because of the manner in which it was extracted from the Janus movement logs.
- Despite being employed in a surveillance mission profile, there appears a *prima facie* case that UAVs could provide higher tier range extension for terrestrial TDDS.
- The metrics that proved particularly useful were 'connectedness', 'phases', 'average path length', 'outage times' and characterisations of the outage times (including peak outage duration).
- These metrics can be used to elucidate characteristics of *ad hoc* networks supporting realistic force deployments and thus assist in defining system requirements.

7. Recommendations

As a consequence, the following recommendations are made:

- Better sources of unit location data than Janus movement logs should be sought.
- Tactical deployments of units within Headline 2000 including the UAV flight profiles should be validated by Army.
- Investigation and endorsement of operationally acceptable peak outage duration and total times should be pursued
- The further work (discussed in para 5) including analysis of a wider set of unit deployment data should be pursued.

8. References

- [1] P.G. Clark and W.D.Blair, "A Summary of the Paper 'Operational Concept Description of the Battlespace Communications System (Land), Final Report' by Frater & Ryan," DSTO Report DSTO-TN-0294, 2000.
- [2] J. Nilsson, A. Hansson, U. Sterner, "The Ability to Provide Services in Military Radio Networks - A Feasibility Study", *Proceeding of IEEE Military Communications Conference (MILCOM2000)*, Los Angeles, California, U.S.A., Oct., 2000.
- [3] Philip Clark and Rowland Dickinson, "Headline 2000 Post Experiment Report - CD Aspects", DSTO Report DSTO-DP-0658, 2001. (Limited Distribution).
- [4] Web Site of the Internet Engineering Task Force Working group Mobile Ad-hoc Networks (manet) located at <http://www.ietf.org/html.charters/manet-charter.html>
- [5] S. Tanenbaum, *Computer Networks*, Third Edition, New Jersey: Prentice Hall International
- [6] I. Chlamtac and S. Kutten, "A spatial-reuse TDMA/FDMA for mobile multi-hop radio networks", *Proceedings IEEE INFOCOM*, Washington, DC, March 1985.
- [7] J. Gronkvist, "Assignment Methods for Spatial Reuse TDMA", *Proceeding of the Sixth Annual International Conference on Mobile Computing and Networking (Mobicom 00)*, Boston, Massachusetts, U.S.A., Aug., 2000.
- [8] A. Ephremides and T.V Truong, "Scheduling Broadcasts in Multihop Radio Networks", *IEEE Transactions on Communications*, **38**, 4, April 1990.
- [9] I. Chlamtac and S.S. Pinter, "Distributed Nodes Organization Algorithm for Channel Access in a Multihop Dynamic Radio Network", *IEEE Transactions on Computers*, **C-36**, 6, June 1987.
- [10] H.L. Bodlaender, T. Kloks, R.B. Tan and J. van Leeuwen "Approximations for λ -Coloring Graphs" Utrecht University Computer Science Technical reports of CS-2000, UU-CS-2000-25 located at <http://www.cs.ruu.nl/pub/RUU/CS/techreps/CS-2000/2000-25.pdf>

Appendix A : Headline 2000 Data Locations

Land Operations Division (LOD) maintains Headline 2000 data and historical information in the Restructuring the Army/Army Experimental Framework (RTA/AEF) knowledge base. This consists of online electronic information, offline electronic information (CD, floppy disk and tape), and hard copy material. The complete experimental results are contained in the RTA Knowledge Repository. This repository is accessible over the Defence Restricted Network at:

<http://yosemite.dsto.defence.gov.au/rtaindex.htm>
using an appropriate web-browser.

For access to offline (including data at 'HE00 data on Doc') and hard copy material please contact the RTA/AEF knowledge base Administrator. At time of publication this is:

Bruce Bennett, LOD, ph 08 8259 6590, email bruce.bennett@dsto.defence.gov.au

Item	Electronic location
Janus data	HE00 data on Doc
Janus runs on CD	by arrangement
Log of Janus runs on CD	... \Trial data\HE00 data\HL00 data lists\HE00 Janus CD+ comments.doc

Appendix B : Graph Theory Measures

B.1 Degree

The degree of the maximum degree node is shown in Figure 15. It is over 40, much higher than experienced in Internet style networks, especially considering the number of nodes involved (generally maximum of 100 or so nodes with periodic merging of 'n' nets into around a subnet of 160). This will drive up the required number of phases and hence reduce individual transmit capacity.




Figure 15 - Maximum Degree over time

B.2 Density

Figure 16 shows the maximum density seen in each time step for the largest connected graph. Note that the density parameter is somewhat non-linear, a small density value can have a large impact on survivability and phase plan. Consider a network of 100 nodes, if all nodes had links to 10 other neighbours then one would consider them to be well provided with alternate routes, yet density is only 0.1.




Figure 16 - Density over time

B.3 Churn




Figure 17 - Average churn per Entity

Figure 17 shows the average churn (as described in 3.2.2) for each entity in the simulation. Churn provides a useful measure as to the dynamism in the network, which should be indicative of the volume of network maintenance data (as opposed to user data) that is required for an entity. The plot does not provide any obvious groupings from which operational entity types (eg helicopters) can be excluded from the tier one network (and provided with other means) based upon detrimental behaviour (i.e large levels of churn).

Appendix C : STDMA as a Graph Problem

In a graph colouring problem, nodes (vertices) or links (edges) of a graph are sought to be coloured following rules particular to the problem. Recall that the rules in our node allocation problem were:

- a node must not transmit in the same phase as any of its neighbours (primary interference)
- two nodes, neighbours to a common node, must not transmit in the same phase (secondary interference)

Another way of expressing this is that any given node and all of its immediate neighbours must have unique phases (so that the given node can have a transmit phase and be able to receive from all of its immediate neighbours). Further, for each neighbour of the given node, none of the subsequent neighbours (i.e. neighbour's neighbours) can use the same phase as the given node (so that each neighbour of the given node will be able to receive from the given node without interference from its own neighbours). Note this second condition does not say anything about uniqueness of phase allocation of these neighbour's neighbours only that they have to be different to the given node. If the neighbour's neighbours share a common neighbour, then their uniqueness is driven by the first rule. Some discussions on node allocation fall into the trap of demanding that all the neighbour's neighbours have unique phase.

The nodes must be coloured (where the colours represent the phase allocated to the node) following the rules⁴ that:

- Nodes at distance 1 from a given node must have a colour differing by at least 1 step in the colour index to the given node.
- Nodes at distance 2 also must have a colour differing by at least 1 step in the colour index to the given node.

Note there is a considerable body of literature on a related problem⁵ where the immediate neighbour's phase differ by 2 steps in the colour index. This embraces the requirement in some FDMA systems for nearby stations to have a greater frequency separation than distant stations to avoid 'adjacent channel' interference.

Using the optimisation approach of seeking to minimise the number of colours, the number of colours required to solve the problem (or 'chromatic power') indicates the number of phases that the network resource needs to be shared across. Assuming each node only receives one phase to transmit, then the larger the number of phases, then proportionately smaller will be the share for each station. In a real network, the graph colouring problem to actually allocate phases (especially in a fast, distributed manner)

⁴ Often referred to as the L(1,1) problem where the parameters show that the immediate neighbours must vary by at least one colour and neighbour's neighbours also by at least one colour.

⁵ Often referred to as the L(2,1) problem

is a key problem. To produce a truly optimal solution is a mathematically intractable problem described as 'NP Complete'. There are, however, sub-optimal solution techniques that are scalable to large numbers of nodes in the network. From an analytical point of view, we are not so much concerned about actually allocating the phases, we are interested in the chromatic power to determine the net capacity of a node to transmit into the network.

In [10] it is shown that the $L(1,1)$ problem is fundamentally the same problem as the traditional graph colouring problem (i.e. $L(1)^6$) of the 'square' of the original graph. The square of a graph is a new graph which has the same vertices (nodes) as the original graph, but where the neighbours were either neighbours (i.e. distance 1) or neighbour's neighbours (i.e. distance 2) in the original graph. The number of colours required for the $L(1)$ solution of a graph is at most the maximum degree+1. If we call the value of the maximum degree of the original graph d , it can be shown that the maximum degree of the square of the graph is at most d^2 . Thus the $L(1,1)$ solution would have no more than d^2+1 colours. However, if the graph is highly meshed (or more correctly, highly meshed around the highest degree node), then the maximum degree of the square of the graph (let us call it $d(G^2)$) will be much less than d^2 . This is because in a highly meshed graph, my neighbour's neighbour will often be a neighbour of mine in the original graph, thus my degree will not change very much when the graph is squared. Density also impacts on how close an estimate of colours is to $d(G^2)+1$. In a fully meshed network the number of colours required for the $L(1)$ problem and the $L(1,1)$ problem become simply $d+1$, but since in a fully meshed network d is actually the number of nodes-1, then the number of colours is the number of nodes.

Regardless, the effort required to square the graph is significant so an easier way is sought. Some work has been carried out by the authors to explore whether a relationship could be discovered between the number of colours for the $L(1,1)$ problem, the maximum degree and density. Results were encouraging, but not sufficiently accurate for our purposes. The reason is suggested to be that density is a network wide measure, whereas what is needed is density around the key nodes that are driving the number of colours. There are difficulties in determining the key nodes and this local density and make it similar in effort to actually allocating the colours.

⁶ The parameters show that only one's immediate neighbours are of concern and they must be different by at least one on the colour index (i.e. not be the same colour)

Appendix D : Linking to University of Adelaide Traffic Model

The Teletraffic Research Centre (TRC) of the University of Adelaide have been tasked (under task LRR01/090) to develop a traffic model from the Headline 2000 data. At this stage the model examines all informed, once per time period, location reporting traffic. The model does not yet include a second tier to bind together the subnets (separated due to terrain effects – see Figure 6).




Figure 18 - Node maximum transmit bit rate required for Situation Awareness data dissemination

Figure 18 is a chart showing the transmission capacity required for the station with the maximum transmission load in bits per second for each time slot under examination. Note that the model employs a 'shortest path' paradigm in determining relaying responsibilities. While this should result in minimising total network load (through minimisation of message relay), it is unlikely to provide for an even traffic distribution. An alternative paradigm may reduce the traffic load on the maximum transmission station but is the subject of further study in LRR01/090. Also note that all nodes were not fully connected in the analysis, so the maximum bit rate is calculated only for

sharing the location data of the subset of nodes in the particular fully connected graph. During the time step interval 1800 to 1900 the nodes did merge into a fully connected network, so this time period provides a good indication as to the traffic load for a complete network (i.e. around 7 kbps).

The TRC provided to the authors the entity identity of this maximum load station, from this, we can determine the subnet that encompasses this station and hence the number of time slots required to meet the STDMA plan for that subnet. Assuming all stations in the subnet have equal access to the medium (this is a constraint that may be removed in future analyses) then the channel bit rate required to sustain the flow of location traffic can be determined (and shown in the figure).




Figure 19 - Channel bit rate requirement

Figure 19 shows the channel bit rate required for each time slot. Of particular interest are periods such as time step 1800 to 1900 where the size of the subnet approaches that required to encompass all operational entities. These periods give a good indication of what channel bit rate is required to carry the location traffic for the entire Headline 2000 force. The value of 350 kbps obtained from this time period is comparable with radio systems that are currently available. For example, the US Army Near Term Digital Radio (NTDR) provides a channel bit rate of 288 kbps and the Enhance Position Locating and Reporting System (EPLRS) radio is expected to reach a channel bit rate of around 500 kbps in the near future. However, the graph only accounts for situational

awareness data and by itself is consuming considerable amounts of channel capacity. Careful generation and dissemination strategies will be needed to minimise all traffic types, especially network maintenance ('router' traffic) and situational awareness traffic. Nevertheless, the reader should note the limitations of the TRC traffic model in respect of load sharing, also, these peak traffic demands (when the subnets merge) occur when the number of phases required peaks and this may be slightly reduced by maintaining links between the two large subnets only via a second tier. Such consideration is a topic for further study in co-operation with the LRR task.

DISTRIBUTION LIST

Ad hoc Networks:
Headline 2000 Communications Analysis

W.D. Blair and A.B. Reynolds

AUSTRALIA

DEFENCE ORGANISATION

Task Sponsor

Director General Command, Control, Communications and Computing

S&T Program

Chief Defence Scientist	}	shared copy
FAS Science Policy		
AS Science Corporate Management		
Director General Science Policy Development		
Counsellor Defence Science, London (Doc Data Sheet)		
Counsellor Defence Science, Washington (Doc Data Sheet)		
Scientific Adviser to MRDC Thailand (Doc Data Sheet)		
Scientific Adviser Joint		
Navy Scientific Adviser (Doc Data Sheet and distribution list only)		

Scientific Adviser - Army
Air Force Scientific Adviser
Director Trials

Electronics and Surveillance Research Laboratory

Chief of Communications Division
Chief of Land Operations Division
Research Leader Military Information Networks
Research Leader Land Weapons Systems
Head Network Architectures Group
Network Architectures Group (Fern Hill Park)
W.D. Blair
A.B. Reynolds
M. Brennan (LOD)

DSTO Library

Library Edinburgh 1 copy
Australian Archives
Library Canberra

Capability Systems Staff

Director General Maritime Development (Doc Data Sheet only)
Director General Land Development
Director General Aerospace Development (Doc Data Sheet only)

Army

Director General Future Land Warfare
Land Warfare Development Centre (7 copies)
Commander
LTCOL D. Webster
LTCOL P. Gates
ABCA National Standardisation Officer, Land Warfare Development Sector,
Puckapunyal (4 copies)
SO (Science), Deployable Joint Force Headquarters (DJFHQ) (L), Enoggera QLD
(Doc Data Sheet only)

Intelligence Program

DGSTA Defence Intelligence Organisation
Manager, Information Centre, Defence Intelligence Organisation

Acquisitions Program

PD JP2072
PD Land75
PD Land125
Land Engineering Agency (2 copies: Attention Mr G. Lampard and Library)

Defence Libraries

Library Manager, DLS-Canberra
Library Manager, DLS - Sydney West (Doc Data Sheet Only)

UNIVERSITIES AND COLLEGES

Australian Defence Force Academy
Library
Head of Aerospace and Mechanical Engineering
Serials Section (M list), Deakin University Library, Geelong, VIC
Hargrave Library, Monash University (Doc Data Sheet only)
Librarian, Flinders University

OTHER ORGANISATIONS

National Library of Australia
NASA (Canberra)
AusInfo
State Library of South Australia

OUTSIDE AUSTRALIA**INTERNATIONAL DEFENCE INFORMATION CENTRES**

US Defense Technical Information Center, 2 copies
UK Defence Research Information Centre, 2 copies
Canada Defence Scientific Information Service, 1 copy
NZ Defence Information Centre, 1 copy

ABSTRACTING AND INFORMATION ORGANISATIONS

Library, Chemical Abstracts Reference Service
Engineering Societies Library, US
Materials Information, Cambridge Scientific Abstracts, US

Documents Librarian, The Center for Research Libraries, US

INFORMATION EXCHANGE AGREEMENT PARTNERS

Acquisitions Unit, Science Reference and Information Service, UK

Library - Exchange Desk, National Institute of Standards and Technology, US

SPARES (copies)

Total number of copies:

66

DEFENCE SCIENCE AND TECHNOLOGY ORGANISATION DOCUMENT CONTROL DATA				1. PRIVACY MARKING/CAVEAT (OF DOCUMENT)	
2. TITLE <i>Ad Hoc Networks: Headline 2000 Communications Analysis</i>			3. SECURITY CLASSIFICATION (FOR UNCLASSIFIED REPORTS THAT ARE LIMITED RELEASE USE (L) NEXT TO DOCUMENT CLASSIFICATION) Document (U) Title (U) Abstract (U)		
4. AUTHOR(S) W.D. Blair and A.B. Reynolds			5. CORPORATE AUTHOR Electronics and Surveillance Research Laboratory PO Box 1500 Edinburgh SA 5111 Australia		
6a. DSTO NUMBER DSTO-TR-1317		6b. AR NUMBER AR-012-358		6c. TYPE OF REPORT Technical Report	
				7. DOCUMENT DATE June 2002	
8. FILE NUMBER E8709/9/12		9. TASK NUMBER JTW99/141		10. TASK SPONSOR DGC4	
				11. NO. OF PAGES 33	
				12. NO. OF REFERENCES 10	
URL on the World Wide http://www.dsto.defence.gov.au/corporate/reports/DSTO-TR-1317.pdf			14. RELEASE AUTHORITY Chief, Communications Division		
15. SECONDARY RELEASE STATEMENT OF THIS DOCUMENT <i>Approved for public release</i>					
OVERSEAS ENQUIRIES OUTSIDE STATED LIMITATIONS SHOULD BE REFERRED THROUGH DOCUMENT EXCHANGE CENTRE, DSTO RESEARCH LIBRARY, EDINBURGH, 5111 SOUTH AUSTRALIA					
16. DELIBERATE ANNOUNCEMENT No Limitations					
17. CASUAL ANNOUNCEMENT Yes					
18. DEFTTEST DESCRIPTORS Tactical Communications, Radio Communications					
19. ABSTRACT This report examines unit location data and terrain from Headline 2000 to investigate communications networks within the manoeuvre units. Within the Enhanced Combat Force timeframe, such units would be supported by a Tactical Data Distribution Sub-system (TDDS). This is envisaged as an <i>ad hoc</i> network changing its topology as units manoeuvre across the battlespace. The report describes approaches that can be applied to Headline 2000 and other data to characterise the nature of the network topology, both statically and as it changes over time, in order to explore the impact on the TDDS network capacity and hence to assist in elucidating consequent requirements of a candidate TDDS.					